IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:	
Sashikanth Chandrasekaran, et al.	Group Art Unit: 2457
Serial No.: 09/872,891	Examiner: Gold, Avi M.
Filed: May 31, 2001	Confirmation No. 3158
For: METHOD AND MECHANISM FOR PREDICTING DATA CONFLICTS AND GENERATING A LOAD DISTRIBUTION PLAN IN A MULTI-NODE SYSTEM	

NOTICE OF APPEAL &
REQUEST FOR PRE-APPEAL BRIEF CONFERENCE

Mail Stop AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In response to the Advisory Action mailed February 17, 2010, Applicant herein submits a Notice of Appeal pursuant to 37 C.F.R. § 41.31(a), and respectfully request for a pre-appeal brief conference.

Claims 1-9, 11, and 13-66 stand rejected under 35 U.S.C. § 102(e) as allegedly being anticipated by U.S. Patent No. 6,691,067 (Ding).

A. Ding does not disclose or suggest *executing a workload first on a single node*, before it is sent to a plurality of nodes for *execution* (i.e., again).

Claim 1 recites executing *the workload* on a single node before *the workload* (i.e., the same workload executed on the single node) *is sent to a plurality of nodes for execution* (Emphasis Added). Claims 14, 32-36, 40, and 54 recite similar limitations. According to the Office Action, column 6, lines 14-32 of Ding allegedly disclose the above limitations. However, the cited passage of Ding actually describes an enterprise management system 180 for monitoring, analyzing, and managing resource usage on "heterogeneous computer systems 150 across the enterprise 100" (c6:24-27). As described in Ding, the computer programs for the

enterprise management system 180 itself may be executed on "one or more computer systems" (c6:14-18). Since Ding specifically describes that the computer programs for the enterprise management system 180 be either run on either "one *or* more computer systems" (i.e., in the alternative), to the extent that the computer programs in Ding is analogized as the claimed workload, Ding clearly does not disclose or suggest executing the workload once at a single node, before it is sent to a plurality of nodes for execution again. There is nothing in Ding that discloses or suggests that after the computer programs for the enterprise management system 180 are executed in a single node, the same computer programs are sent to a plurality of nodes for execution again.

According to the Advisory Action, the claims do not describe that the workload is to be executed once, and then sent to the plurality of nodes for execution again. Applicant respectfully disagrees. The claims clearly recite *executing* the workload on a single node before the workload is sent to a plurality of nodes for *execution* – i.e., again (Emphasis Added). As discussed, Ding does not disclose or suggest these limitations since Ding specifically requires the program be executed either on one *or* more computer systems (not once at a first node, and then it is transmitted to a plurality of nodes for execution again). For at least the foregoing reasons, Applicants submit that the prima facie case of the § 102 rejection has not been established, and request that the § 102 rejection be withdrawn.

B. Ding does not disclose or suggest tracing the execution of the workload to identify a potential data conflict that comprises a potential conflict in the data.

Claim 1 also recites tracing the execution of the workload to identify a potential data conflict, *wherein the potential data conflict comprises a potential conflict in the data* (Emphasis Added). Claims 14, 32-36, 40, and 54 recite similar limitations. According to the Office Action, column 6, lines 28-42 of Ding disclose collecting and monitoring "metric data." However, according to Ding, "a metric is a measurement of a particular system resource. For example, in the preferred embodiment, the enterprise management system 180 collects metrics such as CPU, disk I/O, file system usage, database usage, threads, processes, kernel, registry, logical volumes, and paging" (c6:36-41). Thus, the metric data in Ding is a measure of system resource, and is clearly not a conflict in data (much less, a "potential" conflict in data).

According to the Advisory Action, some metrics such as disk I/O and registry would be subject to conflict, and as such the data that would be collected on those metrics must include

potential conflicts. However, Applicant respectfully notes that there is nothing in Ding that discloses or suggests that the metric data includes any identification of potential data conflict. It is unclear why the metric of Ding that is specifically configured to measure system resource would be purported in the Office Action to also measure data conflict. From the disclosure of Ding, it is clear that the metric is for measuring usage of system resource, which does not require any identification of data conflict, much less potential data conflict. Thus, the § 102 rejection should be withdrawn for these additional reasons.

C. Ding does not disclose or suggest based on a result of the tracing, forming a workload distribution scheme that distributes the workload across the plurality of nodes.

Claim 14 recites *based on* a result of the tracing (i.e., tracing to identify a *potential conflict in the data*), forming a workload distribution scheme that distributes the workload across the plurality of nodes (Emphasis Added). Claims 33 and 34 recite similar limitations. Ding also does not disclose or suggest such limitations. According to page 13 of the Office Action, column 11, lines 11-32, and column 21, line 53 to column 22, line 4 of Ding allegedly disclose the above limitations. However, the cited passages of Ding actually describe altering the "configuration of the enterprise" based on a determined "utilization" (wherein utilization refers to CPU queue lengths, disk queue lengths, amount of memory, etc. – see column 11, lines 49-53). Notably, the resource utilization in Ding is not a potential data conflict. Thus, Ding discloses altering the enterprise configuration *based on utilization*, and clearly does not disclose or suggest forming a workload distribution scheme *based on a result of tracing that identifies potential data conflict*, as described in the claims.

Also, according to the Office Action, the Examiner considers the altering of enterprise configuration in Ding to be the claimed "forming a workload distribution scheme." However, for the sake of argument, even if this were true, Ding still does not disclose or suggest any workload distribution scheme that is formed *based on a result of tracing to identify potential data conflict.* As discussed, the resource utilization in Ding is not potential data conflict. Therefore, forming a workload distribution scheme based on resource utilization does not anticipate forming a workload distribution based on a result of tracing to identify potential data conflict. For these additional reasons, Applicants submit that the prima facie case of the § 102 rejection has not been established for claims 14, 33, and 34, and request that the § 102 rejection be withdrawn for these claims.

D. Ding does not disclose or suggest based on a result of the tracing, optimizing the distribution of the workload across the plurality of nodes.

Claim 32 recites *based on* a result of the tracing (i.e., tracing to identify a *potential conflict in the data*), optimizing the distribution of the workload across the plurality of nodes (Emphasis Added). Claim 35 recites similar limitations. Ding also does not disclose or suggest such limitations. According to page 10 of the Office Action, column 7, lines 6-15, and column 11, lines 11-32 of Ding allegedly disclose "predicting the behavior of the workload across the plurality of nodes." However, Applicants wish to clarify that claims 32 and 35 do not recite "predicting the behavior of the workload across the plurality of nodes." Rather, these claims recite "optimizing the distribution of the workload across the plurality of nodes" based on a result of the tracing, which is not disclosed or suggested in Ding.

Also, according to page 13 of the Office Action, column 11, lines 22-32 and column 21, line 53 to column 22, line 4 of Ding allegedly disclose an alteration of an enterprise, which the Examiner analogized as the claimed "optimizing the distribution of workload across the plurality of nodes." However, for the sake of argument, even if this were true, Ding still does not disclose or suggest any optimization of workload distribution that is *based on a result of tracing to identify potential data conflict*. As discussed, the resource "utilization" in Ding refers to CPU queue lengths, disk queue lengths, amount of memory, etc. – see column 11, lines 49-53. Thus, the resource utilization in Ding is not a potential data conflict. As such, optimizing a workload distribution *based on resource utilization* does not anticipate optimizing a workload distribution that is *based on a result of tracing to identify potential data conflict*.

For these additional reasons, Applicants submit that the prima facie case of the § 102 rejection has not been established for claims 32 and 35, and request that the § 102 rejection be withdrawn for these claims.

Claims 2, 37, 41, 55, 60, and 65

Claim 2 recites that the act of identifying potential data conflicts comprises predicting *how many* data conflicts will occur (Emphasis Added). Claims 37, 41, 55, 60, and 65 recite similar limitations. Ding also does not disclose or suggest such limitations. According to page 3 of the Office Action, column 7, lines 37-67, and column 11, lines 11-32 of Ding allegedly disclose the above limitations. However, the cited passages of Ding actually discloses constructing a so-called "baseline model" from data regarding "performance statistics such as

workload response times, utilization, and throughputs at CPUs, disks, networks, and other elements." However, as similarly discussed, the performance data in Ding does not include any conflict in data. Thus, Ding clearly does not disclose or suggest predicting *how many* data conflicts will occur. In addition, Applicant respectfully notes that the "workload response times, utilization, and throughputs at CPUs, disks, networks, and other elements" in Ding does not disclose or suggest predicting *how many* data conflicts will occur. For these additional reasons, the § 102 rejection should be withdrawn for claims 2, 37, 41, 55, 60, and 65. Applicant notes that the Advisory Action has not considered the above arguments with respect to predicting "how many" data conflict(s) will occur.

Claims 3, 38, 42, 56, 61, and 66

Claim 3 recites that the act of identifying potential data conflicts comprises *predicting types* of data conflicts (Emphasis Added). Claims 38, 42, 56, 61, and 66 recite similar limitations. Ding also does not disclose or suggest such limitations. According to page 4 of the Office Action, column 11, lines 11-54 of Ding allegedly disclose the above limitations. However, as similarly discussed, the cited passage of Ding actually discloses performance metrics, which are not conflict in data. Also, the act of determining performance metrics in Ding is clearly not predicting types of data conflicts. Thus, Ding clearly does not disclose or suggest predicting *types* of data conflicts. For these additional reasons, the § 102 rejection should be withdrawn for claims 3, 38, 42, 56, 61, and 66. Applicant notes that the Advisory Action has not considered the above arguments with respect to predicting "types" of data conflicts.

The Commissioner is authorized to charge any fees due in connection with the filing of this document to Vista IP Law Group's Deposit Account No. **50-1105**, referencing billing number **OID-2000-017-01**.

Respectfully submitted,

DATE: February 24, 2010 By: /Gerald Chan/
Gerald Chan
Registration No. 51,541

VISTA IP LAW GROUP, LLP
1885 Lundy Ave., Suite 108
San Jose, California 95131
Telephone: (408) 321-8663 (Ext. 203)
Facsimile: (408) 877-1662